

# Understanding Windows Storage IO Access in the Age of SSDs

SSDs have brought unique benefits to enterprise storage. Their high speed, low power and low latency are driving emerging (and standard) applications toward new performance thresholds.

SSDs have also introduced a new concept into the storage market: storage devices wear as they are written (versus HDDs, which wear on data access and write). The market saw initial user trepidation, partly because users didn't know their IO patterns and had difficulty estimating them. As a result, many storage architects deployed SSDs with endurance ratings far greater than their workloads needed.

As SSD endurance needs have changed (trending toward lower endurance), proper matching has become imperative – especially in the era of read-centric, lower write endurance quad-level cell (QLC) NAND SSDs like the Micron 5210 ION.

In this paper we will show how a tool built into Windows and another distributed by Micron can help you better understand workload and application storage IO profiles to help match the right SSD to the right deployment.

This paper will show you how to:

- Configure and use Performance Monitor, a system IO monitor tool built into Windows
- Measure SSD-specific wear using Micron's Storage Executive
- Match SSD type and endurance ratings to applications and their workloads

# Background/Problem

SSDs have brought unique benefits to client and enterprise storage. SSDs have also introduced the concept of storage devices that wear as they are written and require a warranted write endurance limit (SSD endurance).

Measuring storage IO profiles for different applications and workloads will enable you to make purchase decisions based on your workload needs. This has become increasingly important with the growth of enterprise SSDs use.

## SSD Endurance Trends

When SSD adoption was just starting (ca. 2007), the idea of a storage device that wears when written was new. Because of this, many system designers initially overestimated the amount of wear their applications would require from SSDs. In the early days of SSDs, 10, 20 or more drive writes per day (DWPD) was normal. It ensured a safety margin.

However, trends now show SSD endurance measured in DWPD is rapidly decreasing. This suggests a better understanding of workload read/write profiles.

Figure 1 shows how DWPD requirements have decreased over time.

## SSD Wear


As SSDs are written and re-written, their wear state (a sort of fuel gauge) is communicated to the host system. SATA SSDs, for example, communicate their degree of wear through SMART.

It is important to note that once an SSD reaches its wear threshold, it transitions into a read only state. The data on the SSD is still available to be read, but that data can't be overwritten, nor can new data be added.



NAND media on SSDs has to be erased before it can be rewritten. This two-step process is called a P/E cycle. Different NAND types support different P/E cycle counts, giving different SSDs different levels of write endurance.

75% Enterprise SSDs Shipped Worldwide (2017): ≤ 1DWPD



**Figure 1: SSD Endurance Trends**  
(Source: Analyst consensus, Forward Insights Datacenter, May 2018)

## Solution

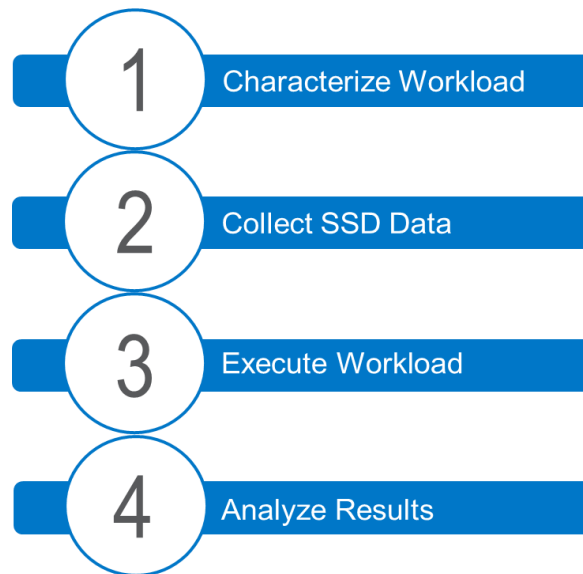
To determine an SSD's suitability for a specific application/workload, one must characterize the IO profile of the workload. This characterization is typically a four-step process, as shown in Figure 2.

The process combines SSD-specific and application specific data to estimate the amount of data written to the SSD (as a function of time) and the workload's read/write ratio.

Many operating systems support integrated tools to help measure the read ratio of data IO on storage devices. When configured to monitor specific OS counters, these tools can keep a running log of storage IO transactions.

Micron supplies SSD-level monitoring software that helps characterize the amount of data written to storage through self-monitoring, analysis and reporting technology (SMART) standard reporting.

This combination of OS and Micron tools can paint a more complete picture of how applications use storage and align this to a specific SSD through a more complete understanding of the workload's storage IO profile.



**Figure 2: Workload Profiling Steps**

### 1 Characterize Workload: Estimate Read/Write Ratio with PerfMon, a Built-In Tool for Windows

Microsoft® Windows® operating systems include **Performance Monitor** (PerfMon), a tool that can be used to glean a better understanding of the IO stream sent to storage devices. This tool enables you to monitor storage IO in real time. (Note that this should be done by an Administrator.)

**Note:** Different versions of Windows may have a slightly different user interface for PerfMon. The essentials for its configuration and use are similar. First create a Data Collector Set, then create and start the new Data Collector.

## Creating a Data Collector Set

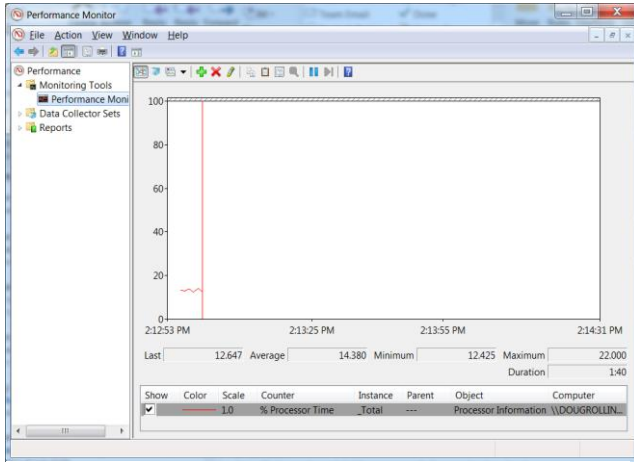


Figure 3: PerfMon UI

1. Launch PerfMon (from the Start Menu, Administrative Tools or via search).

PerfMon uses counters to monitor system activity, including processors, services, network, event tracing and a litany of others. (This paper is focused solely on storage-related counters.)

The next section shows how to add the counters that are most important to storage IO monitoring.

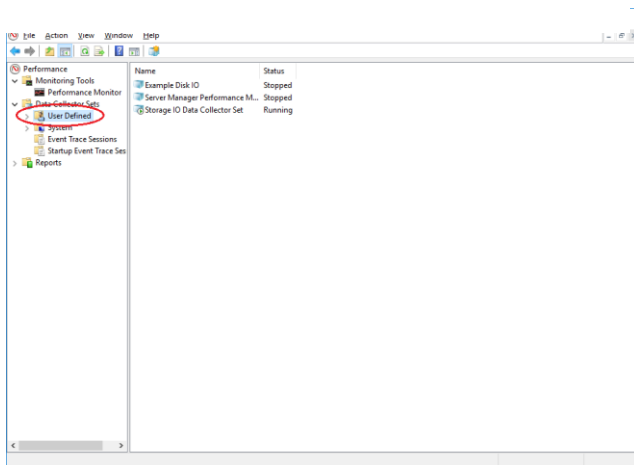


Figure 4: New Data Collector

2. Create a user-defined data collector set (to monitor storage IO).

Right click on User Defined and select Data Collector Set from the context menu.

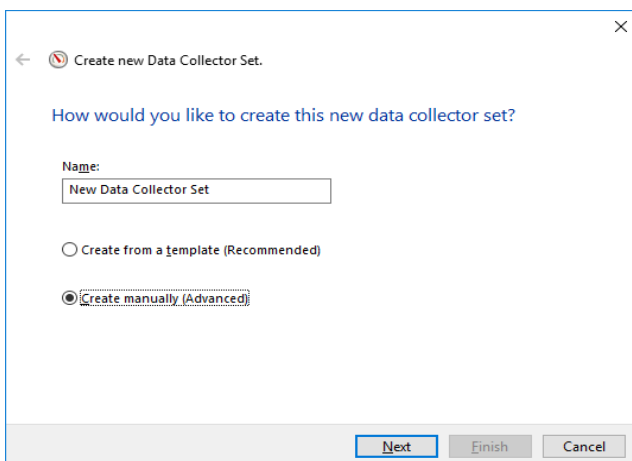


Figure 5: Create Manually

3. Choose how to create data collector set.

In the Name field, change the name if desired (the name is not critical).

Select "Create manually (Advanced)."

This option gives additional choices later in the process. (For this analysis, we'll want to choose the output file format manually.)

Click Next.

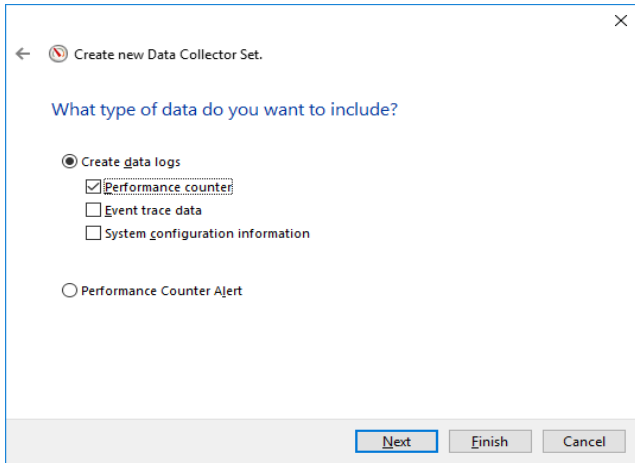


Figure 6: Performance Counter Log Type

4. Choose the type of data log.

Select “Create data logs” and check the box next to “Performance counter.”

Click Next.

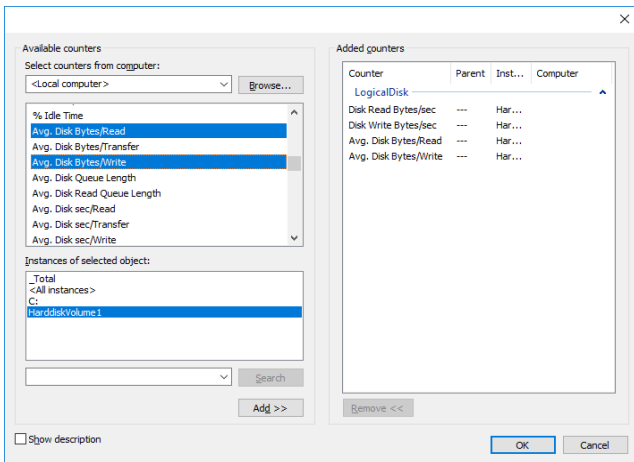


Figure 7: Storage-Related Counters

5. Add counters to the data collector set.

Here, we want to add 4 storage counters.

We want to understand storage IO sizes, as well as storage IO rates (which can be combined with test run time to estimate total bytes read or written).

Click Add.

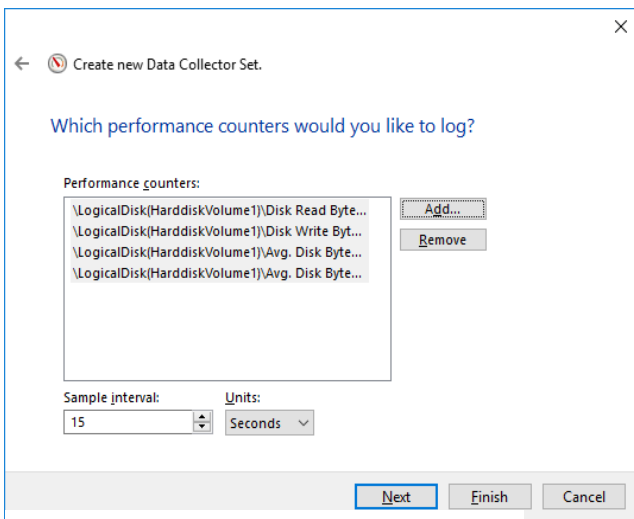


Figure 8: Add Storage-Related Counters

6. Specify counters to add

Expand “Logical Disk” (the storage device to which the workload will be applied). From the list of available Logical Disk counters, choose

- Disk Read Bytes/sec
- Disk Write Bytes/sec
- Avg. Disk Bytes/Read
- Avg. Disk Bytes/Write

The first two counters capture the bytes read/written to storage in one second; the second two counters capture the average read and write transfer sizes.

Select the storage device to monitor and click Add, then click OK, then click Next.

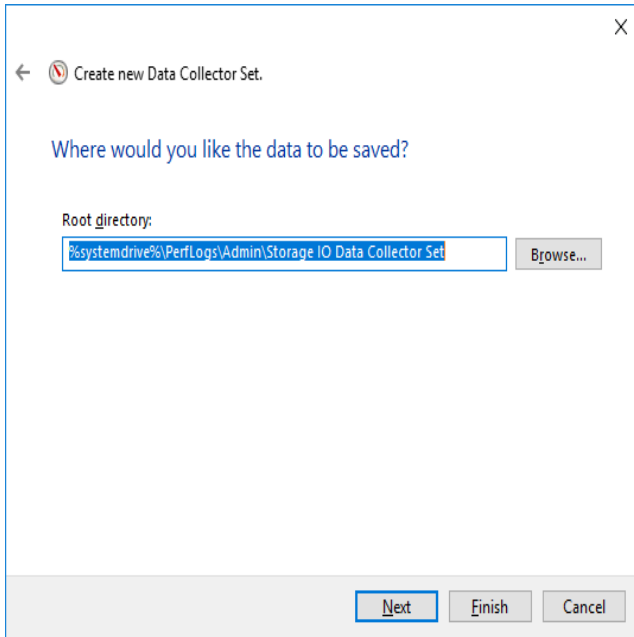


Figure 9: Folder Location

7. Choose the folder location to which you want the data saved.

Click Next and select the folder where you want the new Data Collector Set to be saved.

Click Next, click Save and Close, then click Finish.

The new Data Collector Set is defined. The new Data Collector Set appears under User Defined Data Collector Sets.

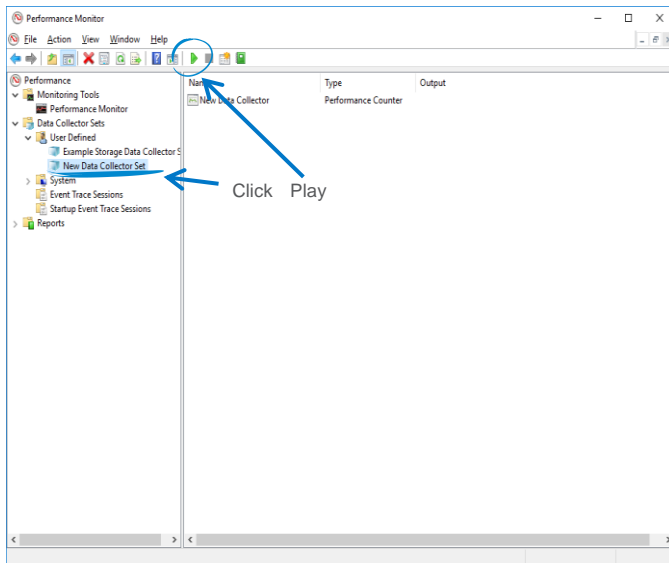


Figure 10: Starting a Data Collector

8. Start the Data Collector.

Click the data collector to highlight, then click the play button. The data collector is now running.

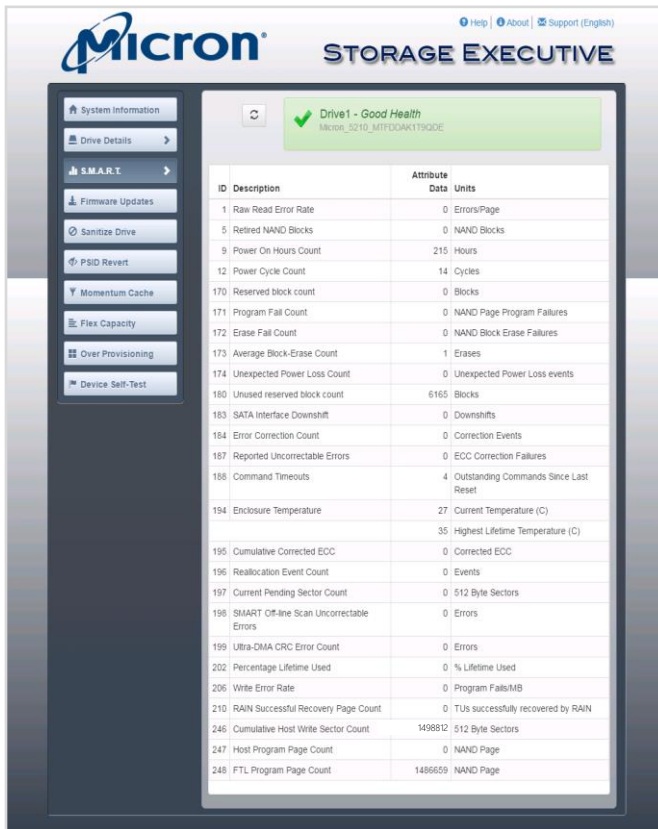


Figure 11: Micron Storage Executive

## Collect SSD Data: Capturing SMART Data

Micron offers a convenient tool called Storage Executive to manage our SSDs and collect data from them (including wear indications expressed via SMART data in the SSD). Storage Executive Storage Executive provides broad functionality, including the ability to:

- View all drives installed in a system and see current drive status and capacity, temperature, firmware version, and driver information
- View SMART attributes
- Update firmware
- Remove all data from a drive by performing a sanitize drive or physical security ID (PSID) revert operation
- Enhance burst performance by enabling the Momentum Cache feature
- Improve drive performance and endurance by allocating over-provisioning capacity
- Increase drive endurance by setting Flex Capacity (supported drives only)
- Perform a drive self-test

This paper uses Micron Storage Executive to monitor how much data is written to the SSD under test by collecting that SSD's SMART data before and after the workload

test.

There are third party tools that can be used to collect SMART data, however while these tools typically collect the data correctly they often misidentify the SMART attributes. If you use a third-party tool, check the Micron SSD datasheet to identify the SMART attribute number associated with expressing the amount of data written not its name. This potential mismatch is only seen in some third-party tools and does not affect the measured results.

Please visit for details on obtaining and using the tool: <https://www.micron.com/products/solid-state-storage/storage-executive-software>

It is also important to note that SMART wear data is cumulative, it does not reset (or return to a prior value) therefore one needs to collect SMART data before and after running a workload test.

The difference between ID 246 (Cumulative Host Write Sector Count) before and after the test will indicate the amount of wear applied during the test.

Figure 13 shows an example output from Storage Executive for a Micron 5210 ION SSD. SMART ID 246 shows the amount of data written to this SSD in 512-byte increments (sectors).

To determine the total amount written to this SSD during any workload test:

1. Record the starting value of ID 246
2. Run the complete workload test (record the time to complete the test)

3. Record the ending value of ID 246

The difference between the values of ID 246 in step 3 and step 1 equals the total number of 512-byte sectors written to the SSD. (It may be easier to express this value in KB or MB through arithmetic conversion for subsequent analysis.)

Note that these calculations will be relative to a single SSD; hence, they may differ from PerfMon values noted earlier (this is expected).

Ensure you've recorded SMART ID 246 value before continuing to the next section.

### 3 Execute Workload

Workload execution is application-specific. Ensure you have captured the starting SMART ID 246 value and that PerfMon (with the correct Collector Set) is running. Record the total duration of the workload when execution completes. Once the workload completes, stop the Data Collector Set and record the ending SMART ID 246 value.

### 4 Analyze Results

Storage Read/Write Ratio: PerfMon Data

Locate the .CSV file that contains the PerfMon logs for the workload completed. The default location is %systemdrive%\PerfLogs\Admin\

Time	Disk Read Write Bytes/sec	Disk Write Bytes/sec	%Read	%Write
START	32785.99594	655.7199188	98	2
	12309.02695	0	100	0
	83580.96896	1671.619379	98	2
	15898.47984	0	100	0
	12824.14445	0	100	0
	17938.43318	179.3843318	99	1
	722060.9874	0	100	0
	82564.63518	0	100	0
	8099.388815	0	100	0
	16411.28004	164.1128004	99	1
END	32317.4592	323.174592	99	1

Figure 12: Example PerfMon Data

#### Estimating Read/Write Ratio: Disk Bytes/sec (using Disk Read Bytes/sec, Disk Write Bytes/sec)

The counters shown in Figure 12 capture the total number of bytes written to the monitored storage device and read from the same device over a period of one second. If the interval is set for a longer amount of time, PerfMon averages the values. Rounding counters to integer values may also ease analysis.



Estimating read/write ratios can help quantify a workload as read-centric, write-centric or mixed use. This in turn can help identify an SSD whose endurance characteristics are well-matched to the workload. While estimating workload read/write ratios using PerfMon data helps target a specific SSD type, this data shows the IO patterns being sent to the storage device and does not factor IO concatenation or bifurcation that may take place at different layers of the storage IO stack and therefore may not directly translate directly into the IO pattern seen by the SSD.

### Example Data Written to SSD (as seen by the SSD): SMART Data

We can monitor the exact amount of data written to the SSD using its wear reporting mechanism. (Note that this example uses SMART; other SSDs may use different reporting methods).

When SSDs report the amount of data written to them (in this example via SMART; other SSDs may differ), one can determine the amount of data written by capturing the SSD’s SMART data at the start and end of the test. In this example, we use SMART ID 246 – the number of 512-byte sectors the host wrote.

Table 1 shows an example. We subtract the ID 246 starting value from its ending value and convert the result (512-byte sectors) into MB. The result is the total MB written to the monitored storage device during the workload test. Combining this data with the test run time yields the test workload’s data write rate.

SMART ID 246	
Starting Value	209,715,200
Ending Value	209,971,200
#512 Byte Sectors Written	256,000
MB Written to SSD	about 125

**Table 1: SMART Calculations**

### Example Workload Analysis

One can combine the SSD choice targeting data from PerfMon, SMART data from an example drive and total test run time for better insight on SSD choice. This example analysis uses a hypothetical SSD with a warranted TBW of 870TB over 5 years. (Note that this is an example only and actual results may differ. Estimating a workload-SSD match should be done using the actual SSD being considered and the workload applied.)

#### PerfMon Data

Suppose our example workload’s PerfMon data is consistent with the example data in Figure 12 – about 99% read. This suggests that our example workload is a good candidate for a read-centric SSD. (Different PerfMon data may suggest other SSD types; for example, when PerfMon data shows a heavy write emphasis as might be seen in a data logging workload, a write-centric SSD may be optimal. When the data shows mixed read and write IO as might be seen in a transaction processing workload, a mixed-use SSD may be optimal.) PerfMon data can help target initial SSD selection.

#### SMART Data

Once we’ve targeted an SSD type (in this example PerfMon data suggests we target a read-centric SSD), we can estimate whether the rated endurance of the read-centric SSD we are considering is sufficient to support a 5-year replacement schedule.

Suppose we run the test workload continuously on the example SSD and use a 30-second test interval and collect the SMART data seen in Table 1. This SMART indicates that over the course of a 30-second test, the workload wrote about 125MB to the SSD (about 4.16 MB/s).

Assuming the workload's write rate is constant and that this 30-second sample reflects the workload IO profile over time, we can extrapolate the total amount of data that would be written to the SSD during its 5-year warranty period and compare that to the SSD's rated endurance of 870TB.

$$(4.16\text{MB/sec}) * (5 \text{ years}) = \sim 656 \text{ TB written}$$

Since our example SSD is rated for 870TB over 5 years, we see that the measured data suggests our application would write less than this SSD's 5-year TBW (SMART data suggests this application would write about 656TB over 5 years).

## Benefits

Understanding a workload's read/write ratio can help identify potential SSDs which may benefit the workload. For example, a read-centric workload would excel with read-centric SSDs and would not benefit from the extra endurance of higher write endurance, higher-cost SSDs. Combining the read/write ratio and measured data written over time helps validate whether an SSD's rated endurance is sufficient.

Profiling storage IO with measured data is the best way to ensure an optimal match between workload and SSD. While there is a litany of conventional wisdom describing "typical" or "historical" application and workload IO profiles, many of these are configuration-dependent.

Matching appropriate applications and workloads to SSDs has become more important with the introduction of quad-level cell (QLC) SSDs. These SSDs are designed for read-centric workloads (and provide better longevity when writing large IO sizes). When SSDs first introduced the concept of storage devices that wear into the storage market, this concept was new – but today easily available tools enable us to more precisely understand our storage IO needs.

## Conclusion

Built-in and Micron-supplied tools combined with workload execution timing can help characterize application and workload storage IO profiles. This, in turn, enhances understanding of data placement and a more precise match between SSD and workload.

**micron.com**

All information herein is provided on an "AS IS" basis without warranties of any kind, including any implied warranties, warranties of merchantability or warranties of fitness for a particular purpose. No hardware, software or system can provide absolute security and protection of data under all conditions. Micron assumes no liability for lost, stolen or corrupted data arising from the use of any Micron product, including those products that incorporate any of the mentioned security features. Products are warranted only to meet Micron's production data sheet specifications. Products, programs and specifications are subject to change without notice. Dates are estimates only.

©2019 Micron Technology, Inc. All rights reserved. Micron, the Micron logo, and all other Micron trademarks are the property of Micron Technology, Inc. All other trademarks are the property of their respective owners. Rev. A 02/19 CCM004-676576390-11247